



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04C 12/66, H04J 3/14		A1	(11) International Publication Number: WO 00/35130
			(43) International Publication Date: 15 June 2000 (15.06.00)
(21) International Application Number: PCT/US98/27774		(74) Agents: OGAWA, Richard, T. et al.; Townsend and Townsend and Crew LLP, Two Embarcadero Center, 8th floor, San Francisco, CA 94111-3834 (US).	
(22) International Filing Date: 4 February 1999 (04.02.99)			
(30) Priority Data: 60/110,976 4 December 1998 (04.12.98) US 09/205,063 4 December 1998 (04.12.98) US		(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).	
(63) Related by Continuation (CON) or Continuation-in-Part (CIP) to Earlier Applications US 60/110,976 (CIP) Filed on 4 December 1998 (04.12.98) US 09/205,063 (CIP) Filed on 4 December 1998 (04.12.98)			
(71) Applicant (for all designated States except US): UKIAH SOFTWARE, INC. [US/US]; Suite 210, 2155 S. Bascom Avenue, Campbell, CA 95008 (US).		Published With international search report.	
(72) Inventors; and (75) Inventors/Applicants (for US only): RAKOSHITZ, Gregory [US/US]; 675 Florales Drive, Palo Alto, CA 94306 (US). VAID, Aseem [IN/US]; 3019 Payne Avenue, San Jose, CA 95128 (US). PUTTA, Sanjeev [IN/US]; Apt. E-2, 2530 S. Bascom Avenue, Campbell, CA 95008 (US).			
(54) Title: DIRECTORY ENABLED POLICY MANAGEMENT TOOL FOR INTELLIGENT TRAFFIC MANAGEMENT			
(57) Abstract			
<p>A method and system (100) for monitoring or profiling quality of service within one or more information sources in a network of computers. The method includes a step of providing a network of computers, each being coupled to each other to form a local area network. The network of computers has a firewall server (110) coupled to the network of computers and a distributed traffic management tool coupled to the firewall server. The method also includes implementing traffic monitoring or profiling of incoming and outgoing information from one of the information sources.</p>			
<p>SCHEMATIC OF QOS SYSTEM DEPLOYED FOR INTERNET ACCESS AND A PRIVATE WAN</p>			

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

DIRECTORY ENABLED POLICY MANAGEMENT TOOL FOR INTELLIGENT TRAFFIC MANAGEMENT

5

10

15

BACKGROUND OF THE INVENTION

The present invention relates to communication or telecommunication. More particularly, the present invention provides a technique, including a method and system, for monitoring and allocating bandwidth on a plurality of locations or nodes in a telecommunication network at, for example, a firewall access point and other positions. As merely an example, the present invention is implemented on a wide area network of computers or workstations such as the Internet. But it would be recognized that the present invention has a much broader range of applicability including local area networks, a combination of wide and local area networks, and the like.

Telecommunication techniques have been around for numerous years. In the early days, people such as the American Indians communicated to each other over long distances using "smoke signals." Smoke signals were generally used to transfer visual information from one geographical location to be observed at another geographical location. Since smoke signals could only be seen over a limited range of geographical distances, they were soon replaced by a communication technique known as telegraph.

Telegraph generally transferred information from one geographical location to another geographical location using electrical signals in the form of "dots" and "dashes" over transmission lines. An example of commonly used electrical signals is Morse code. Telegraph has been, for the most part, replaced by telephone. The telephone was 5 invented by Alexander Graham Bell in the 1800s to transmit and send voice information using electrical analog signals over a telephone line, or more commonly a single twisted pair copper line. Most industrialized countries today rely heavily upon telephone to facilitate communication between businesses and people, in general.

In the 1990s, another significant development in the telecommunication 10 industry occurred. People began communicating to each other by way of computers, which are coupled to the telephone lines or telephone network. These computers or workstations coupled to each other can transmit many types of information from one geographical location to another geographical location. This information can be in the form of voice, video, and data, which have been commonly termed as "multimedia." 15 Information transmitted over the Internet or Internet "traffic" has increased dramatically in recent years. In fact, the increased traffic has caused congestion, which leads to problems in responsiveness and throughput. This congestion is similar to the congestion of automobiles on a freeway, such as those in Silicon Valley from the recent "boom" in high technology companies, including companies specializing in telecommunication. As 20 a result, individual users, businesses, and others have been spending more time waiting for information, and less time on productive activities. For example, a typical user of the Internet may spend a great deal of time attempting to view selected sites, which are commonly referred to as "Websites," on the Internet. Additionally, information being sent from one site to another through electronic mail, which is termed "e-mail," may not 25 reach its destination in a timely or adequate manner. In effect, quality of service or Quality of Service ("QoS") of the Internet has decreased to the point where some messages are being read at some time significantly beyond the time the messages were sent.

Quality of Service is often measured by responsiveness, including the 30 amount of time spent waiting for images, texts, and other data to be transferred, and by throughput of data across the Internet, and the like. Other aspects may be application specific, for example, jitter, quality of playback, quality of data transferred across the Internet, and the like. Three main sources of data latency include: the lack of bandwidth

at the user (or receiving) end, the general congestion of Internet, and the lack of bandwidth at the source (or sending) end.

A solution to decreasing data latency includes increasing the bandwidth of the user. This is typically accomplished by upgrading the network link, for example by 5 upgrading a modem or network connection. For example, the network link may be upgraded to X2 modems, 56K modems, ADSL or DMT modems, ISDN service and modems, cable TV service and modems, and the like. Drawbacks to these solutions include that they typically require additional network service; they also require additional hardware and/or software, and further they require both the sender and receiver to both 10 agree on using the same hardware and/or software. Although one user may have a much faster line or faster modem, another user may still rely on the same 1,200 kbaud modem. So, the speed at which information moves from one location to another location is often determined by the slowest information which is being transferred over the network. Accordingly, users of faster technology are basically going nowhere, or "running" 15 nowhere fast, as is commonly stated in the network industry.

From the above, it is seen that a technique for improving the use of a wide area network is highly desirable.

SUMMARY OF THE INVENTION

20 The present invention relates to a technique, including a method and system, for providing more quality to telecommunication services. More particularly, the present invention relates to quality of service management using a novel traffic monitoring technique, which is distributed over a network. The present monitoring technique is predominantly software based, but is not limited to such software in some 25 embodiments. The present invention also provides a management tool for allocating bandwidth, as well as other features.

In a specific embodiment, the present invention provides a system with a novel graphical user interface for monitoring a flow of information coupled to a network of computers. The flow of information can come from a variety of location or nodes such 30 as a firewall, a server, a wide area network, a local area network, a client, and other information sources. The user interface is provided on a display. The display has at least a first portion and a second portion, where the first portion displays a graphical chart representing the flow of information, which comes from one of many locations on the

network. The second portion displays text information describing aspects of the flow of information. The combination of the first portion and the second portion describes the information being profiled. The display also has prompts in graphical or text form or outputs the source of the flow of information, where the source can be one of a plurality 5 of nodes such as a server, a firewall, a wide area network, a local area network, a client, and other information sources. The present invention can be distributed over a network by way of one or more agents.

In an alternative specific embodiment, the present invention provides a novel computer network system having a real-time bandwidth profiling tool. The real-10 time bandwidth profiling tool has a graphical user interface on a monitor or display. The graphical user interface includes at least a first portion and a second portion. The first portion displays a graphical chart representing the flow of information from at least one information source. The second portion displays text information describing the flow of information. The combination of the first portion and the second portion describes the 15 information being profiled. Additionally, the graphical user interface has a portion that outputs a graphical representation including text or illustration of the source itself. The flow of information can be from a variety of sources, such as those described above as well as others, to provide a distributed profiling tool.

In still an alternative embodiment, the present invention provides a novel 20 bandwidth profiling tool. The present bandwidth profiling tool includes a variety of computer codes to form computer software or a computer program, which is stored in computer memory. The program includes a first code that is directed to measuring a data rate for a flow of information from an incoming source, which is coupled to a node from one of a plurality of sources such as a network of computers, for example. The program 25 also has a second code that is directed to categorizing the data rate from the flow of information based upon at least one of a plurality of traffic classes and a third code that is directed to outputting a visual representation of the data rate in graphical form on a display. A fourth code is used to direct the outputting of a text representation of the one of the plurality of traffic classes on the display. A fifth code is used to display the origin 30 of the flow of information. The origin of the flow of information can be one of a plurality of nodes from a firewall, a server, a local area network, and wide area network, and others. The present invention has a variety of other codes to perform the methods described herein, and outside the present specification.

Numerous advantages are achieved by way of the present invention over pre-existing or conventional techniques. In a specific embodiment, the present invention provides a single point or a single region to manage telecommunication traffic including directory services and bandwidth management. Additionally, in some, if not all 5 embodiments, the present invention can be implemented at a single point of access such as a computer terminal or firewall, for example. Furthermore, the present invention can be predominately software based and can be implemented into a pre-existing system by way of a relatively simple installation process. Moreover, the present invention provides more valued applications and users with a more reliable and faster service. Less critical 10 applications and users are provided with a service level that is appropriate for them in some embodiments. In most embodiments, available bandwidth in a system is fairly shared between equally prioritized users (e.g., no user can monopolize or "hog" the system). Still further, link efficiency improves due to overall congestion avoidance in most cases. Moreover, the present invention implements its traffic management 15 technique using a simple and easy to use "rule" based technique. Still further, the present invention has tools that are distributed at one or more locations on the network to monitor traffic on an enterprise level rather than a single point or node on the network. Accordingly, the present invention provides an "end to end" full cycle traffic management 20 program. Depending upon the embodiment, one or more of these advantages can be present. These and other advantages are described throughout the present specification, and more particularly below.

Further understanding of the nature and advantages of the invention may be realized by reference to the remaining portions of the specification, drawings, and attached documents.

25

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a simplified diagram of a system according to an embodiment of the present invention;

Fig. 2 is a simplified block diagram of system architecture according to an 30 embodiment of the present invention;

Fig. 3 is a simplified diagram of a traffic management cycle according to an embodiment of the present invention;

Figs. 4-7 are simplified diagrams of systems according to various embodiments of the present invention;

Fig. 8 is a simplified flow diagram of a rule-based control method according to the present invention;

5 Figs. 9-15 are simplified representations of graphical user interfaces for monitoring traffic according to the present invention; and

Figs. 16-19 are simplified diagrams of a distributed bandwidth management system according to embodiments of the present invention

10 DESCRIPTION OF SPECIFIC EMBODIMENTS

An embodiment of the present provides integrated network service policies for firewall platforms, as well as other platforms or gateways. Specifically, the present invention provides network or firewall administrators with the ability to implement policy-based schema for security and resource management on firewall platforms. In a specific embodiment, resource management includes Network Quality of Service (QoS) or "bandwidth" management techniques. In an exemplary embodiment, the present invention provides tools for monitoring traffic for bandwidth management, as well as other functions.

20 Network QoS occurs by managing the resources that serve network application traffic, for example. This typically includes the following resources: link bandwidth, application server bandwidth (CPU), and buffer space on generally all nodes (end-points, routers and gateways). Typically, data through-put is limited by the speed of Internet access links and by the server CPU capacity, and response time is determined by the number of hops in a route, physical length of the route, and extent of congestion in the route. There are various other factors that may affect QoS, such as the behavior of TCP/IP, severe congestion anywhere in the route, prioritization of traffic along the route, etc. To a network administrator, embodiments of the present invention provide discrimination of different traffic types and provide methods for enforcement of traffic flow by management to the above resources.

25 30

DEFINITIONS

In the present invention, it may assist the reader to understand some of the terms described herein. These terms have been briefly described below. These terms are merely examples and should not unduly limit the scope of the claims herein.

5 1. Traffic Management: A set of techniques or mechanisms including policies that can be applied in a network to manage limited network resources such as bandwidth and the like. These techniques are intended to improve overall network performance and efficiency. They are also intended to provide for more predictability and orderliness in the event of network congestion. The techniques should also isolate 10 faults and provide visibility into performance problems. Additionally, they should meet the diverse user and application requirements as per an organization's business goals. Furthermore, traffic management is intended to increase the "good put" traffic, based on the economic value and prevent the abuse of network resources.

15 2. Quality Of Service (QoS): The concept of Quality of Service (QoS) has been analyzed and discussed for a number of years in the networking industry, and was previously associated mostly with ATM technology. In a more generic sense, QoS describes the performance specifications that an application requires from the underlying infrastructure. Otherwise, the application will not run satisfactorily. Some 20 applications are designed to run in a best-effort mode and can adapt to available bandwidth. Others are extremely sensitive to delays. Still others can produce large bursts in traffic which affects other applications while providing little perceptible improvements to the end-user. QoS specifications are closely associated with the expectations and perceptions of end-users and the organization they are part of.

25 3. Bandwidth: Bandwidth usually refers to maximum available bit rate for a specific application. In a specific embodiment, synchronous, interactive, and real-time applications, which are bandwidth-sensitive, can require minimum bandwidth guarantees, and can require sustained and burst-scale bit-rates. On the other hand, 30 network administrators may want to limit bandwidth taken by non-productive traffic such as push technologies like PointCast and others. Even though bandwidth may be allocated for specified applications, it does not mean that these applications may be using that

bandwidth. Therefore, a good policy should be to enforce when there is competition and demand.

4. Latency: Latency generally refers to the delay experienced by a 5 packet from the source to destination. Latency requirements are typically specified as mean-delay and worst case delay in some cases. Real-time audio/video applications such as, for example, DNS, HTTP, and TELNET are delay sensitive. Delay is a result of 10 propagation delay, due to physical medium and queuing at intermediate nodes such as routers, gateways, or even servers. A certain portion of the delay can be controlled by how the queues are serviced at the intermediate nodes, and by controlling congestion at bottleneck points. Some examples of delay measures are packet round-trip delay and connection response time.

5. Jitter: Jitter generally refers to variation in delay (e.g., that is, the 15 delay is not constant for all packets of a given flow) for a particular application. Real-time applications require a worst case jitter. Applications such as real-audio and video do some advanced buffering to overcome any variation in packet delays - the amount of buffering is determined by the expected jitter.

20 6. Packet Loss: Packet loss is a loss in a packet or a portion of packets that is generally caused by failure of network elements (e.g., routers, servers) to forward or deliver packets. Packet loss is usually an indication of severe congestion, overload of an element, or element failure (e.g., if a server is down). Even if the packet was not dropped but just delayed, protocols and applications can assume it was lost. 25 Packet loss can cause application timeouts, loss of quality or retransmitted packets. Packet loss is usually specified as a rate (e.g., a real-time video application cannot tolerate loss of more than one packet for every 10 packets sent). Indirect results of packet loss may also be measured (e.g., connection retries or data retransmits).

30 7. Guarantees: An extreme example of a guarantee is to partition bandwidth so that it is not available to other entities. Guarantee also means a share of the resource, e.g., minimum bandwidth or maximum latency.

8. Best-effort: Best-efforts describes a service on best-effort basis but makes no guarantees.

9. Limits: Specific physical or theoretical limitation on a resource 5 such as bandwidth. Resource utilization or admission is limited under certain conditions.

10. Priority: Level of importance for a specific user, application, or data. Create a priority scheme among different entities so that contention is resolved or service is provided.

10

11. Traffic Profiling: Profiling is intended to be defined as cumulative details of traffic flows for each active client, server, or application without application of any rules. This includes bandwidth, response time, and failure related statistics. Profiling is intended to provide long term cumulative snapshots of traffic for capacity planning or 15 setting traffic rules.

20 The above definitions are merely intended to assist the reader in understanding some of the terms described herein. They are not intended, in any manner, to limit the scope of the claims. One of ordinary skill in the art would recognize other variations, modifications, and alternatives.

SYSTEM OVERVIEW

Fig. 1 illustrates a simplified system 100 according to an embodiment of the present invention. The system 100 is merely an illustration and should not limit the 25 scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. The present invention, which includes a bandwidth management tool, can be embodied as a TrafficWare™ firewall server 110 from Ukiah Software, Inc, but can be others. The bandwidth management tool, which sits in the firewall, can monitor and/or control information at a single node or more than one 30 node (i.e., distributed) in the network system. System 100 typically includes a file server 120, and a plurality of computers 130-150, coupled to a local area network (LAN) 160, and other elements. Firewall server 110 includes a typical connection to a wide area

network (WAN) 170 and to a remote LAN 180 (such as an Intranet) and a typical network connection 190 to the Internet 200. Attached to Internet 200 are Web servers 210 and other computers 220.

As illustrated, computers such as computer 130, 140, and 210 5 communicate using any one or multiple application layer protocols such as Telnet, file transfer protocol (FTP), Hypertext transmission protocol (HTTP), and the like. Further, communication across WAN 170 and across network connection 190 implements transport layer protocols such as transmission control protocol (TCP), universal data protocol (UDP), and the like. LAN 160 and LAN 180 are preferably based upon network 10 protocols such as Internet protocol (IP), IPX from Novell, AppleTalk, and the like. As shown in Fig. 1, network connection 190 may be accomplished using T1, ISDN, Dial-up, and other hardware connections. Computers 120-150 and 210-220 may be any suitable make or model of computer that can be coupled to a network. The system can also include a variety of other elements such as bridges, routers, and the like.

15 In an alternative specific embodiment, the present invention may be applied to a system with various links accessed in servicing a browser request at a remote web server. In this embodiment, a client could be dialing in via a 28.8 kbit dial up modem to a local Internet service provider (ISP), where the ISP may be connected to the Internet by a T1 link. A web server may be on a 10 BMS Ethernet LAN, which is 20 connected to another ISP via a 56 K frame relay. The web server's ISP may be connected to its carrier via a T3 line. The client ISP carrier and the server ISP carrier may both be connected by an ATM backbone or the like. Because of this asymmetry in this embodiment, any traffic management solution should take into account these variations including traffic speed and data format described above. Moreover, simply upgrading the 25 capacity of a link, in the access path, may not offer a viable solution. This present embodiment may have the bandwidth requirements shown by way of Table 1, for example.

Table 1: Bandwidth Requirements

Users	Bandwidth	Service Offered
Internet developers, individuals, international locations where bandwidth is expensive	28.8 to 56 Kbps	Dial-up services, ISDN
Small to medium-sized organizations with moderate Internet usage	56 Kbps to 1.5 Mbps	Fractional TI, frame relay
Medium sized organizations with many moderate users, smaller organizations requiring huge amounts of bandwidth	1.5 Mbps	Dedicated TI circuit
Standard bandwidth for Ethernet-based LANs	10 Mbps	Ethernet, token ring (4 Mbps or 16 Mbps)
Bandwidth usage for large organizations or Internet backbones	45 Mbps	Dedicated T3 circuit
Huge bandwidth LAN backbone usage for medium to large organizations (hundreds or thousands of users)	100 to 1.000 Mbps	Fast Ethernet, gigabit Ethernet

As shown above, there exist a large number of diverse applications and protocols that are widely used and have their own performance requirements. For example, applications such as mail (e.g., SMTP) and news (e.g., NNTP) are not

5 interactive and are therefore not sensitive to delay. On the other hand, applications such as real-time conferencing are extremely sensitive to delay but not to packet loss.

Applications such as TELNET and DNS do not utilize significant bandwidth, but are sensitive to delay and loss. Conversely, applications such as FTP consume a great deal of bandwidth but are not that sensitive to delay. Generally, network applications can be

10 categorized as:

1. Interactive (e.g., delay sensitive) versus non-interactive (e.g., delay tolerant);
2. Bandwidth intensive (bulk data) versus non-bandwidth intensive;
3. Bursty versus non-bursty.

15 These categories are merely illustrative and should not limit the scope of the claims herein. Additionally, some application requirements are dependent on the context of use and the nature of data being accessed. Such applications can be described as being nominally interactive or nominally bandwidth intense. This means the 20 description applies to many but not all the situations in which they are used.

As merely an example, Table 2 provides some illustrations for these categories.

Application Class	Examples
Low-bandwidth, delay sensitive, highly interactive	DNS, PING, TELNET, CHAT, COLLABORATION
High bandwidth, delay sensitive	Real-time audio and video
High Bandwidth, nominally interactive	Web service requests, file downloads

Non-interactive	Mail and news
-----------------	---------------

Table 2: Application Spectrum

As shown in Table 2, low-bandwidth, delay sensitive, and highly interactive applications include, among others, DNS, PING, TELNET, CHAT,

5 COLLABORATION. High bandwidth and delay sensitive applications including at least real-time audio and video. Additional applications for high bandwidth and nominally interactive, or non-interactive have also been shown. Again, these applications are merely provided for illustration and should not limit the scope of the claims herein.

The present invention can also be used with a number of various files.

10 For example, a number of common applications, such as FTP and HTTP, can handle a wide variety of files. The file types being transferred and downloaded place different demands on the underlying infrastructure. Index and HTML files take up limited bandwidth but have very mundane contents. On the other hand, GIF, JPEG and MPEG, RA and AVI files take up a lot more bandwidth but provide a rich multimedia experience
15 to the end-user. In fact, push technologies such as PointCast basically download rich-multimedia bandwidth-intensive files.

The present invention can also be used with a variety of user requirements. For example, networks are facing an explosion in the number of (inter) networked applications and data accessible through them. Network resources are increasingly being
20 used for a wide variety of purposes, ranging from business critical to personal. This means that policies must ensure that scarce resources (e.g., Internet bandwidth) are utilized with the goal of maximizing the returns to the organization. These benefits can come from direct revenue generating activities or from improved productivity (or reduced loss of productivity). As shown in Table 3, for example, at a mythical company called
25 "Shebang Software Inc." the highest bandwidth priority has been allocated to technical support. However, there is no hard and fast rule. As with security policies, decisions should be consistent with the needs of the organization.

Table 3: Shebang Software User Priorities

Users	Application Class	Reasons
Technical support	Mission critical	Needs most bandwidth to deal with customers who need assistance
Sales and marketing	Critical	Needs bandwidth to deal with potential customers. Answer inquires, make quotes, transmit multimedia presentations
Upper management and middle management, administrative	Casual	Needs bandwidth to perform tasks necessary to run the business
Development and manufacturing	Personal	Needs bandwidth to send e-mail, subscribe to Push technologies

The present invention takes into account, in one or more embodiments, the factors which are described specifically above. Although the above has been generally described in terms of a specific type of information, other types of information on a network can also be used with the present invention. Additionally, the present invention

has been described in general to a specific system. For instance, the present bandwidth management tool can be applied at a network's Internet access link. Alternatively, the present tool can be applied to a private WAN link to a remote corporate site or an access to a server farm (e.g., a group of servers located in a special part of the network close to an access link, e.g., in a web hosting environment). Alternatively, the present invention can be applied to key servers (e.g., database/web server) within an organization servicing internal and/or external users. Furthermore, the present bandwidth management tool can be applied to any combination of the above or the like. Still further, the tool can be distributed in one or more locations or nodes in the network, e.g., LAN, WAN.

Fig. 2 is a simplified block diagram 200 of details of system architecture according to an embodiment of the present invention. The block diagram is merely an illustration and should not limit the scope of the claims herein. The architecture includes a variety of layers that each interface to each other as depicted by the layers. The system includes a network layer 211, which interfaces to incoming and outgoing information to the network. The network can be one of a variety including, among others, Ethernet and Token Ring. A physical layer 209 is disposed above the network layer 211. The physical layer can be personal computers, which are commonly called PCs, or network interface computers, which are commonly called NCs, or alternatively workstations. As merely an example, a personal computer can be an IBM PC compatible computer having a '586-class based microprocessor, such a Pentium \square from Intel Corporation, but is not limited to such a computer or processor. An operating system ("OS") is used on the computer such as WindowsNT \square from Microsoft Corporation, but can also be other OSs. The system is also coupled to a graphical user interface ("GUI") 201 and is coupled to directory services such as, for example, LDAP, but can be others. A detailed discussion of directory services is described in U.S. Application Serial Nos. _____, (Attorney Docket Nos. 18430-000110, 18430-000120, 18430-000230) which are commonly assigned, and hereby incorporated by reference for all purposes.

Directory services 224 and GUI 201 couple to an application programming interface ("API") 223. The API is coupled to a traffic management or bandwidth management tool 208 with at least three modules, including a policy engine module 231, a FAST module 229, and a FAIR module 227, which will be discussed in more detail below, but is not limited to these modules. The bandwidth management tool 208 can be predominantly software based and is substantially free from any significant hardware or

software changes in the network. In a preferred embodiment, the bandwidth management tool 208 can be loaded onto a server without any changes to hardware. In an alternative preferred embodiment, the tool can install, configure, and operate on a conventional IBM compatible PC running and operating system such as, for example, Windows NT, but can be others. The tool can be deployed at any appropriate point in the network data path. 5 The tool can also be stand-alone at the WAN access point (e.g., behind the Internet access router or behind a firewall), with a conventional firewall or with an NT based proxy/caching server or application server (e.g., a Web server).

Tool 208 performs incoming and/or outgoing management of information 10 over the network of computers. In a specific embodiment, traffic management tool 208 performs inbound and outbound monitoring and control of flows by application, source address, destination address, URL, time of day, day of week, day of month, and other variations. In a specific embodiment, tool 208 also monitors, controls, and produces reports and alarms, which can enhance a whole spectrum of traffic monitoring and control 15 activities ranging from bandwidth/latency control to capacity planning.

In a specific embodiment, the bandwidth management tool adapts to "real" changes on any pre-existing networking system. For example, network infrastructure management involves a continuous process of monitoring, reporting, and deploying changes to match network growth or changing needs in a growing office, for example. 20 These changes exist at various levels and time scales. As merely examples, the network changes can be to enforce a QoS Policy for a critical service, add WAN bandwidth, segment the network, upgrade a router, choose a guaranteed service level for a web site (e.g., user's own web site), or notify "Mr. Hog" (i.e., a user occupying too much bandwidth) that he should schedule his large personal downloads at more prudent times 25 such as late at night, for example.

In a specific embodiment, the system architecture has applications or tools that distribute itself over the network, which allow the present tool to monitor one or more nodes on the network. In one aspect, the present tool can be disposed at any source of information such as a router, server, a firewall, a bridge, a local area network, a wide 30 area network, a client, and other information sources. Further details of the distributed bandwidth management product is shown by way of the Figs. below.

The bandwidth management tool can employ these changes using, for example, the process shown in Fig. 3. This process is merely an illustration and should not limit the scope of the claims herein. As shown, Fig. 3 is a simplified diagram 300 of a traffic management cycle according to an embodiment of the present invention. The traffic management cycle is depicted as a continuous cycle, which includes a monitoring phase 301, a creating/applying policy phase 303, and a reporting/alarming phase 305, but is not limited to these cycles. That is, these cycles can be separated or combined depending upon the application. By way of this cycle, the tool can adapt to any changes to the networking system according to the present invention.

In an aspect of the present invention, the present tool can monitor and control activities at various times, e.g., seconds, days, weeks, months, years. Some details with regard to these control activities are shown below under the headings.

1. Second to second

The tool provides second to second time scale monitoring and control of incoming and outgoing traffic over the network. As merely an example, the tool ensures that critical or more important traffic gets a right of way during traffic bursts and provides bandwidth enforcement. Multiple users of the network at a specific time can cause the traffic burst. Alternatively, multiple sessions on the network at a specific time can cause the traffic burst. Once the traffic burst is detected, the tool has a control device, which provides bandwidth enforcement to ensure that the more important traffic gets through the network.

2. Day to day

The tool provides day to day time scale monitoring and control of incoming and outgoing traffic over the network. As merely an example, the tool manages time of day congestion, and responds to intermittent problems or perceived problems. The tool generally deals with problems or limitations that are very specific and isolated to particular users or particular services at particular times that need to be tracked down quickly.

3. Week to week

The tool provides week to week time scale monitoring and control of incoming and outgoing traffic over the network. The tool analyzes traffic usage performance patterns, what services or hosts are active on the network, and troubleshoots chronic problems. In particular, the tool looks at aggregates, such as a particular segment 5 of the network, and compares Websites or compares groups of users for usage of bandwidth and frequency of usage.

4. Longer term activities

The tool provides long term time scale monitoring and control of incoming 10 and outgoing traffic over the network. The tool implements changes in organizational priorities, in billing. The tool also provides service for new applications as they are introduced, and provides for capacity planning for network resources. The present tool can also be used with network stress testing tools to obtain detailed analysis of flows and traffic behavior with/without policy enforcement before a new application is deployed to 15 change the network infrastructure.

Based upon the above description, the present tool can be used to monitor and control incoming and outgoing traffic over a variety of time frequencies. The time frequencies include second by second, day to day, or long term, and combinations 20 thereof, depending upon the application. Of course, the time frequency used depends upon the particular network and applications.

Figs. 4-7 are simplified diagrams of systems according to various 25 embodiments of the present invention. These diagrams are merely illustrations and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, alternatives, and modifications. These systems show various deployment scenarios according to the present invention.

1. Internet Service Provider (ISP)

Fig. 4 is a simplified diagram 400 of the present tool in an ISP 30 environment according to the present invention. The diagram 400 includes a variety of elements such as an ISP LAN 401, which is coupled to network elements including a remote access concentrator 403, a web server 417, an FTP server 415, a router 413, a news server 411, and others. The tool 405 is coupled between the ISP LAN and router

407, which is connected to the Internet 409. In this embodiment, the ISP is providing a number of services to its customers and the present tool sits by the Internet link and manages inbound and outbound traffic.

In a specific embodiment, the system architecture has applications or tools 5 that distribute itself over the network, which allow the present tool to monitor one or more nodes on the network. In one aspect, the present tool can be disposed at any source of information such as a router, server, a firewall, a bridge, a local area network, a wide area network, a client, and other information sources. As merely an example, the present tool can be implemented at any location that is identified by reference letter B, but can 10 also be at other locations.

2. Web Hosting Deployment

Fig. 5 is a simplified diagram 500 of the present tool in a web hosting environment according to the present invention. The diagram 500 includes a variety of 15 elements such as a LAN backbone 501, which is coupled to network elements including web servers 503, 511, 513, and others. The present tool 505 is coupled between LAN 501 and router 507, which is connected to the Internet 509. In the present embodiment, the tool is being used to manage inbound and outbound traffic between some Websites and the Internet. In a specific embodiment, most of the data being transmitted is multimedia-based, but is not limited as such data. 20

In a specific embodiment, the system architecture has applications or tools that distribute itself over the network, which allow the present tool to monitor one or more nodes on the network. In one aspect, the present tool can be disposed at any source of information such as a router, server, a firewall, a bridge, a local area network, a 25 wide area network, a client and other information sources. As merely an example, the present tool can be implemented at any location that is identified by reference letter B, but can also be at other locations.

3. End-User Deployment

Fig. 6 is a simplified diagram 600 of the present tool in a campus 30 environment according to the present invention. The diagram 600 includes a variety of features such as a campus network 601, which is coupled to network elements such as a desktop PC 603, a UNIX computer 617, an NT Server 615, a web server 613, directory

services 611, and others. A bandwidth management tool 605 is coupled between campus network 601 and router 607, which is coupled to Internet 609. In this embodiment, a LAN or WAN supports a number of different setups and configurations, which are compete for bandwidth to access the Internet. The present tool acts as an arbitrator for 5 implementing rules, enforcing policies, and setting admissions for classes, as well as perform other acts.

In a specific embodiment, the system architecture has applications or tools that distribute itself over the network, which allow the present tool to monitor one or more nodes on the network. In one aspect, the present tool can be disposed at any source 10 of information such as a router, server, a firewall, a bridge, a local area network, a wide area network, a client, and other infomation sources. As merely an example, the present tool can be implemented at any location that is identified by reference letter B, but can also be at other locations.

15 4. Private WAN

Fig. 7 is a simplified diagram 700 of the present tool deployed for a large corporation that has an Intranet as well as an Internet. The diagram 700 includes a variety of elements or "children" such as a connection to Frankfurt 715, a connection to London 713, a connection to Hong Kong 717, and a connection to Paris 719. Each connection or 20 child includes a router 705A, E, D, C, and the present tool 703A, E, D, C, which is coupled between the router and the hub ("HQ"). In a WAN-based environment, for example, HQ 701 is the hub that handles a number of independent systems (e.g., Frankfurt, London, Hong Kong, Paris), which can be LAN-based. In this embodiment, the present tool 703B also sits by the Internet 711 and is used to allocate bandwidth 25 between the competing children, e.g., Frankfurt, London, Hong Kong, Paris. Router 705B is coupled between tool 703B and Internet 711.

In a specific embodiment, the system architecture has applications or tools that distribute itself over the network, which allow the present tool to monitor one or more nodes on the network. In one aspect, the present tool can be disposed at any source 30 of information such as a router, server, a firewall, a bridge, a local area network, a wide area network, a client, or other information sources. As merely an example, the present tool can be implemented at any location that is identified by reference letter B, but can also be at other locations.

Although the above descriptions have been made in terms of deploying the present tool in selected environments, the present tool can also be deployed in other environments. For example, the present tool can be deployed in any combination of the above. Alternatively, the present tool can be deployed in any portion of the above environments. Of course, the type of environment used by the present tool depends highly upon the application.

In a specific embodiment, the tool provides an easy to use interface or graphical user interface ("GUI") for performance monitoring and profiling (e.g., accounting). Profiling can be based on active services, clients and servers, among other parameters. Additionally, profiling of the network can be started as soon as the tool is installed into the server of the network. Accordingly, the tool provides immediate accounting and service measurement on a variety of QoS measures.

In a specific embodiment, the present tool generally uses two mechanisms to implement efficient traffic monitoring and traffic control. These mechanisms include processes performed by the FAST module and the FAIR module, which are shown in Fig. 2, for example. Additionally, the present tool uses a policy engine module 231, which oversees the FAST module 229 and the FAIR module 227. Some details of these modules are described as follows.

20 1. FAST Module (Flow Analysis and Session Tagging)

The FAST module generally provides for monitoring of incoming and outgoing information to and from the network or link. Flow Analysis and Session Tagging ("FAST") implements rich, application level traffic classification, and measurement. This operation is accomplished without introducing slow data paths to minimize latency and maximize overall throughput of traffic through the tool management engine. As shown in the Fig., the FAST module provides for classification 203 of information such as parameters 213 including application, presentation, session, transport, and network. The FAST module also provides for measurement 219 of various parameters. The FAST module is coupled to the API.

30

2. FAIR Module (Flow Analysis and Intelligent Regulation)

The FAIR module generally implements traffic control and manages bandwidth of incoming and outgoing information to and from the network or link. Flow

Analysis and Intelligent Regulation ("FAIR") implements traffic control based on a combination of flow control and queuing algorithms. FAIR's objective provides inbound and outbound traffic management for meaningful time intervals, reducing the load on packet classifiers and packet schedulers. The FAIR module controls 205 incoming and 5 outgoing information to and from the network. Additionally, the FAIR module controls 205 by parameters 215 such as class, session, burst, packet, and others. The FAIR module also controls time 217 of allocating bandwidth for these parameters. The FAIR module is coupled to the API.

10 3. Policy Engine Module

The policy engine module 231 oversees the FAST and FAIR modules. The engine module also interfaces with the API. In an embodiment, the policy engine module includes a security policy 201, a traffic policy 202, and other policies 221. The security policy provides parameters for securing the present tool. The traffic policy 15 defines specific limitations or parameters for the traffic.

Some definitions about the various modules have been described above. These definitions are not intended to be limiting. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Additionally, the modules described are generally provided in terms of computer software. Computer software can 20 be used to program and implement these modules, as well as others. The modules can be combined or even separated, depending upon the applications. Functionality of the modules can also be combined with hardware or the like. In a specific embodiment, the present modules are implemented on an WindowsNTTM operating system, which has been developed by Microsoft Corporation. Of course, other operating systems can also 25 be used. Accordingly, the present modules are not intended to be limiting in any manner.

In an embodiment, the present tool can be configured based upon at least the following components - traffic classes, traffic policies, traffic rules, and traffic entities. Some information about these components are described below.

30 1. Traffic Classes

The present tool identifies data flows at a network site based on traffic classes. A traffic class is any combination of the following, but is not limited to these:

IP address, sub-net, network, net group, or range of source or destination;
URL of the sender or group of URLs;
Service (e.g., HTTP, FTP) or groups of services;
FTP and HTTP, file types can be selected as well;
5 Time of day, day of week/month; and
Inbound and outbound information.

As shown above, traffic classes are directional. Traffic classes configured for inbound traffic are managed separately from traffic classes configured for outbound traffic. For example, the present tool may decide to guarantee a minimum bandwidth to critical traffic so that it is not affected by congestion from large downloads. Additionally, the present tool may want to monitor Push traffic for a while and then choose to limit it if it is perceived as a problem. Traffic classes can also be for measurement only or for control and measurement in some embodiments. These are merely examples and should 15 not limit the scope of the claims herein.

2. Traffic Policies

Traffic policies are generally mechanisms used to control the traffic behavior of specific classes. In an embodiment, the present tool can configure policy 20 properties which provide, for example:

Bandwidth guarantees - granting classes a minimum bandwidth in the presence of congestion or competition;

25 Bandwidth limits - establishing a limit on the total bandwidth used by the class;

Setting priorities - establishing a priority order for bandwidth limiting or servicing traffic from a class. (That is, high priority classes are serviced first and are 30 affected the least during contention for bandwidth. Lower priority classes are serviced in order of priority and may be more affected by congestion or contention);

Admission control- establishing conditions under which a new network session or service request is admitted or not admitted. (This kind of policy establishes a broad bandwidth control or service quality for sessions already admitted).

5 As shown, the present invention provides policies such as bandwidth guarantees, bandwidth limits, setting priorities, admission control, and others. It may assist the reader in understanding some of the terms used in the policies by drawing an analogy with a geographical highway for automobiles. For example, bandwidth relates to how fast one can go (e.g., fast or slow lane) once a user has entered the stream of traffic 10 on the highway. That is, the physical limit for speed in the specific lane chosen. Priority is analogous to how quickly the user is able to enter the highway and move into a designated lane, and how often the user may have to temporarily give way to other vehicles during the drive. Admission control is analogous to the metered lights at the entrance of the freeway where one is made to wait under certain conditions. Of course, 15 depending upon the applications other analogies can be used to explain the policies. Additionally, the policies are merely examples and should not limit the scope of the claims herein.

3. Traffic Rules

20 A rule generally includes a traffic class and a policy associated with the class. A class can have several policies that apply at different time intervals. 'Rule' is also used to refer to the policy or to a specific row in the present tool user interface. The present tool user interface is described in, for example, U.S. Application No. _____ (Attorney Docket No. 18430-000300, commonly assigned, which is hereby incorporated 25 by reference for all purposes.)

4. Traffic Entities

The present tool refers to entities in at least two different contexts: defining traffic classes and viewing traffic profiles. For example, a network entity 30 generally refers to an IP address, host, sub-net, IP net, IP range, URL or a group of other network entities. A service entity refers to a single service or a group of services. A native entity is referred to in viewing traffic profiles. No rule setting or configuration is required to monitor these entities. When the present tool is installed, it begins to profile

traffic based upon detected services, clients, or servers, all of which are called native entities.

5. Guidelines for Developing Traffic Policies

5 The present invention provides some guidelines for developing traffic policies. For example, to develop meaningful and effective traffic policies, the present tool may need to understand and take into account one or more of the following:

10 •The kind of business being performed by the user over the Internet. If the user is an ISP, the user may need to develop a business/pricing model that leverages the features of the present tool. If the user is managing corporate access to the Internet, the user may want to identify any business critical services being provided over the Internet

15 •The priority of clients, servers and URLs hosted in the user's network or servers access over the Internet. This can be organized as business critical, casual and personal.

20 •The properties of different applications being used, whether they utilize lots of bandwidth or not. The user may also need to account for the type of files commonly download by users or from the Web site.

•Measure and analyze traffic using the present tool's profiles.

Additionally, monitoring of selected entities (e.g., users, services) may also be useful.

25 In a further embodiment, the present tool provides some general guidelines of some commonly used applications. These guidelines should be used in conjunction with business driven priorities, traffic profiling, and selective real-time monitoring to establish an effective traffic policy. Selected guidelines are defined as follows, but are not limited to these.

30

•Delay-sensitive low bandwidth applications, such as TELNET and DNS, are controlled best by setting a high priority policy. The present tool can give the highest

priority to all network control traffic, such as QoS signaling, session establishment, domain lookup and routing protocols.

- Streaming multimedia applications, such as Real Audio/Video and
- 5 Vxtreme, can hog allot of bandwidth but are also delay and bandwidth sensitive. If they are not critical, they are controlled best by setting a high priority and a policy to limit admission of sessions so that bandwidth use is capped but admitted sessions have a reasonable quality.

10 Push technologies, such as PointCast and Marimba, download large files, are not delay or bandwidth sensitive and usually not business critical. They are best controlled by a limiting bandwidth policy and a low priority.

- Bulk-data non-interactive applications, such as SMTP and NNTP, should

15 be guaranteed a small bandwidth minimum so that they are not totally squeezed out by congestion or control policies.

- Bulk-download, nominally interactive applications, such as FTP or some

20 HTTP downloads, are commonly used in a variety of situations, ranging from critical to casual. Differentiating various types of usage in this case can usually be made only on the basis of file types and/or source or destination addresses. In this case, a small minimum can be guaranteed for more important use.

- In bulk-download applications (e.g., file size > 20 K Bytes), overall

25 congestion and burstiness can be controlled by slightly limiting this traffic, even if it is just a little below the total available bandwidth (e.g., 90%). The present tool can provide smoothing controls on this traffic without impacting overall perceptible performance for these downloads. This is particularly useful at lower link speeds (128 K and below).

30

- Mission critical applications, such as Lotus Notes, Oracle SQLNet, and LDAP, are controlled best by setting a high priority with a guaranteed bandwidth minimum.

The above provides some guidelines for commonly used applications according to the present invention. Using the above guidelines, the present tool can effectively allocate bandwidth on a network, for example. Again, the above guidelines are merely examples and should not limit the scope of the claims herein.

5 In a specific embodiment, the present tool provides a comprehensive, flexible, rule-based paradigm for implementing traffic control, as illustrated by a simplified flow diagram 800 of Fig. 8. This flow diagram 800 is merely an illustration and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Before explaining the flow
10 diagram, it may assist the reader by reviewing some general terms used herein.

These terms include, among others, "rules" and "classes" and "policies." Rules can be created for very specific groups of flows or more general groups of flows, which are commonly all the stuff that transmits to and from a link to a gateway point. Groups of flows are also referred to as traffic classes, but are not limited to such classes.
15 Classes also can be defined by source, destination, application, file types, URLs, and other features. Policies can be specified to control traffic flows in terms of overall bandwidth guarantees, bandwidth limits, priority of service, how individual sessions within a class are serviced or admitted, and other aspects. The present tool also has intelligent policy validation that prevents users from defining any contradictory or
20 ambiguous rules. Policy validation is generally a higher level check used by way of the present method.

The present method occurs at start, which is step 801, for example. In general, a flow of information or data or packets of information enter a gateway point, where the present tool sits. The present method classifies (step 803) the flow of
25 information. Groups of flows can be referred to as traffic classes, but are not limited to such classes. Classes also can be defined by source, destination, application, file types, URLs, and other features. Other examples of classes were previously noted, but are not limited to these classes. In general, step 803 classifies the flow of information received into one of a plurality of predetermined classes.

30 The present tool measures parameters for each of the classes in step 805, which were received, for example. These parameters are based upon the policy or rule, which may be applied in a later step. As merely an example, parameters include the class itself, file sizes, and other information, which can be used by the policy or rule to apply

the policy or rule to improve the quality of service for the network. After measuring the parameters, the present method applies a time stamp (step 807) on the parameters to correlate the class of information received to a time, for example.

A step of determining whether to apply a policy occurs in the next step 5 809. For example, if the class and the time (and the link state in some embodiments) meet predetermined settings, the policy is applied to the class in step 811 through branch 810. Alternatively, if one of the elements including the class, the time, or the link state do not meet the predetermined settings, the policy does not apply and the process continues to measure parameters through branch 808. Alternatively, the process continues to 10 measure parameters through branch 821 after the policy is applied to the flow of information for the class.

Depending upon the application, the policy is used to improve the quality of service of the network by performing at least one of a number of functions for the class of information from the flow. These functions include, among others, bandwidth 15 guarantees, bandwidth limits, setting priorities, admission control. The present process can also halt or stop as shown in step 815. The steps occur, in part, by way of the modules, which were previously described, but can also occur using other techniques including a combination of hardware and software, for example. These sequence of steps are merely illustrative and should not limit the scope of the claims herein. One of 20 ordinary skill in the art would recognize other modifications, alternatives, and variations. -

In a preferred embodiment, the present invention uses a variety of graphical user interfaces for profiling and monitoring traffic. Figs. 9A-15 are simplified representations of graphical user interfaces for monitoring traffic according to the present invention. These representations are merely illustrative and should not limit the scope of 25 the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives.

Fig. 9A is a simplified flow diagram 950 of a profiling method according to the present invention. Profiling or monitoring traffic can occur using one of a plurality of user interfaces or graphical user interfaces. The present invention provides a profiles 30 tab 953, which can be selected using a mouse or keyboard interface. The present method begins with a start step, which is step 951. Upon selecting a profiles tab 953, one of a plurality of tabs is prompted. These tabs represent services 957, server 959, and client 961. These tabs display relevant traffic statistics by every active service, server and

client, respectively. By selecting one of the tabs, the present tool sorts data or information in ascending order by clicking on any header (e.g., Kb Transferred), as illustrated by Fig. 9 for a service tab 900. Other functions that can be performed using one of the profiles and the graphical user interface include:

5

- Click the Refresh button, all data is updated from the profiling engine.
- Click the Reset button 907, clears all the respective data from the profiling engine.

10

- Click the Save As 909 button to save the respective data to a log file. The data is saved as tab-separated text.

Each of the present user interfaces also includes function keys 901 and a 15 tool bar 903. Upon selecting the profiles tab, a profiles light or display indication illuminates 911. As shown, the main profiles tab also includes tabs for services 913, server 915, and client 917. Additional features of the various tabs including the services tab, the server tab, and the client tab are described below and refer to Figs. 9, 10, and 11, respectively, but are not limited to these descriptions.

20

1. Services Tab

Fig. 9 is a simplified diagram 900 of a representation of a graphical user interface for a services tab according to the present invention. In particular, the dialog box displays cumulative traffic statistics for selected applications. The services tab, 25 which can be selected by default, provides the following information:

Service Name

This field 919 shows what services (e.g., All Services, FTP, HTTP, 30 SMTP, POP3, SSL) the network uses. Summary statistics for all services (e.g., inbound or outbound) are also shown. Traffic from services that are not recognized by the present tool are indicated as 'Others'.

Direction

30

This field 919 indicates whether the service is inbound or outbound.

Note: Inbound and Outbound refer to the direction of data flow, not the request.

5

Kb Transferred

This field 923 shows the amount of data transferred in inbound or outbound direction. As shown, the amount of data can be in kilobits transferred. Additionally, the amount of data can be referred to as a percentage of all services.

10

Connect Response Time

This field 925 indicates an average time to establish a session. The connect response time is in milliseconds, but is not limited to this time. The minimum and maximum connect response time is also shown in parenthesis.

15

Request Response Time

This field 927 indicates an average response time for an application request. The request response time is in milliseconds, but is not limited to this time. The minimum and maximum request response time is also shown in parenthesis.

20

Note: This measure is application specific and does not apply to all services. For example, for HTTP, it is the time taken by a URL to start sending data after a request for a file was made by a Web browser.

25

Total Sessions

This field (not shown) indicates the total number of sessions established for this service.

Retries

30

This field (not shown) indicates the percentage of connect requests that needed to be retried. Retries can result from network congestion, packets dropped in the network or server congestion.

Server Aborts

This field (not shown) indicates the percentage of sessions aborted by the server.

5

Time

This field (not shown) indicates the last time the service was active.

2. Server Tab

Fig. 10 is a simplified diagram 1000 of a representation of a graphical user interface for a server tab according to the present invention. Upon selecting or clicking the server tab 915, screen 1000 appears. The dialog box displays cumulative traffic statistics for every active server. The server tab provides the following information, but is not limited to such information:

15

Server

This field 1001 shows the server host name, URL or IP address. Summary statistics for all servers are also shown.

Note:

20

- In one aspect of the invention, the present tool can profile up to 256 servers. Subsequent traffic from new servers are indicated as 'Others'.

- Host names can also be displayed in some embodiments.

25

Kb Transferred

This field 1003 shows the amount of data transferred from the server. As shown, the amount of data can be in kilobits transferred. Additionally, the amount of data can be referred to as a percentage of all services.

30

Round Trip Time

This field 1005 indicates an average round trip delay for packets sent to the server. The round trip time is in milliseconds, but is not limited to this time. The minimum and maximum round trip time is also shown in parenthesis.

Connect Response Time

This field 1007 indicates an average time to establish a session with the server. The connect response time is in milliseconds, but is not limited to this time. The 5 minimum and maximum connect response time is also shown in parenthesis.

Total Sessions

This field 1009 indicates the total number of sessions established to the server.

10

Retries

This field (not shown) indicates the percentage of connect requests that needed to be retried. Retries can result from network congestion, packets dropped in the network or server congestion.

15

Server Aborts

This field (not shown) indicates the percentage of sessions aborted by the server.

20

Access Speed

This field (not shown) indicates the bottleneck speed for the route between the present tool as a host and a server.

Data Retransmits

25

This field (not shown) indicates the percentage of data packets that were retransmitted by the server.

Time

30

This field (not shown) indicates the last time data was received from the server.

3. Client Tab

Fig. 11 is a simplified diagram 1100 of a representation of a graphical user interface for a client tab according to the present invention. When the client tab 917 is selected or is clicked using a user interface, screen 1100 appears. The dialog box displays the cumulative traffic statistics for the clients. The client tab provides the following 5 information, but is not limited to such information:

Client

This field 1101 shows the client host name or IP address. Summary statistics for all clients are also shown.

5

Note: The present tool can profile up to 256 clients in some embodiments. Subsequent traffic from the clients are indicated as 'Others'.

Kb Transferred

10

This field 1103 shows the amount of data transferred to the client. As shown, the amount of data can be in kilobits transferred. Additionally, the amount of data can be referred to as a percentage of all services.

Round Trip Time

15

This field 1105 indicates an average round trip delay for packets from this client. The round trip time is in milliseconds, but is not limited to this time. The minimum and maximum round trip time is also shown in parenthesis.

Connect Response Time

20

This field 1105 indicates the average time to establish a session from the client. The connect response time is in milliseconds, but is not limited to this time. The minimum and maximum connect response time is also shown in parenthesis.

Total Sessions

25

This field 1109 indicates the total number of sessions established from the client.

Retries

30

This field (not shown) indicates the percentage of connect requests that needed to be retried. Retries can result from network congestion, packets dropped in the network or server congestion.

Server Aborts

This field (not shown) indicates the percentage of sessions aborted by the server.

Time

5 This field (not shown) indicates the last time the client received data through the link used by the present tool.

10 The present invention provides the aforementioned tool for profiling a variety of information from a flow of information at a communication link. The tool has an easy to use graphical user interface, which can sort information by at least services, client, or server, depending upon the application. The illustrations shown are merely used 15 as examples and should not limit the scope of the claims herein.

15 In a specific embodiment, the present invention with graphical user interface begins profiling upon installation. In particular, the present tool is installed onto a server to automatically start profiling traffic in inbound and outbound directions without any further configuration. The present tool can be stopped and restarted manually from a user interface. While the present tool is stopped, profiling is interrupted temporarily.

20 The present invention provides additional easy to use graphical tools to monitor and profile traffic. In one aspect, the present invention takes advantage of a Windows NT™ Performance Monitor to monitor traffic for any measurement or control - rule that is created. In another aspect, the present invention can launch the Performance Monitor from the 'Administrative Tools' Program group and select counters for monitoring incoming and outgoing traffic from a link.

25 Fig. 12 is a simplified graphical user interface 1200 to launch a performance monitoring tool according to the present invention. This interface is merely an illustration and should not limit the scope of the claims herein. A method for launching the present tool occurs, in part, by selecting or clicking on the performance monitor tab 1201. The display shows available traffic classes 1201 (e.g., FTP, HTTP, PointCast), which have been defined in the traffic policy. Note that a traffic class is not a 30 rule. There may be more than one rule that belongs to the same traffic class. Traffic classes are created when rules are edited. A traffic class is defined by at least a source, destination, and service properties. The display includes a group of option buttons 1207

titled monitor, which allows a user to specify whether the user wants to monitor bandwidth consumption 1209, connect time 1211, or connect retries 1213 for the selected classes. A prompt box 1215 above the option buttons 1207 provides a brief explanation of the selected option. A Launch button 1205 launches the performance monitor too. To 5 launch the present performance monitor tool:

1. Select one or more traffic classes 1203 in the list.
2. Choose monitor by clicking on an appropriate option button (e.g., 10 bandwidth consumption, response time, failures) 1207 in the monitor group.
3. Push launch button 1205.

As merely an example, Fig. 13 is a simplified graphical user display 1300 15 for bandwidth consumption according to the present invention. As shown, the Fig. is an example of Class Bandwidth 1305 monitoring for a few services 1307 such as FTP, HTTP, etc. over a 56 Kbit Internet link. The vertical axis 1302 illustrates a bandwidth scale from "0" to "56.0" kbits and the horizontal axis represents time 1306. The plurality of line plots 1304 each represent one of the services 1307, which are each color coded 20 1301 for easy reading by the user. The display also includes an object 1309 and a computer 1311, which is being used to monitor the traffic. Accordingly, the present display includes a graphical portion 1310 and a text portion 1320. The graphical portion includes the plurality of plots representing the services for bandwidth consumption as functions of time. The text portion is in the form of a legend, but can also include other 25 information.

The illustration in the above Fig. is merely an example and should not limit the scope of the claims. Although the present example has been described in terms of bandwidth consumption, the present performance monitor tool can also be used to monitor a variety of other parameters, as discussed above. These other parameters 30 include, among others, connect time, or connect retries for the selected classes. Furthermore, the present tool has other types of charts such as a bar chart, a pie chart, and the like. Of course, the parameter being profiled and monitored depends upon the application.

In an alternative embodiment, the present invention provides a user interface for modifying the plots or charts, such as the one previously described, as well as others. Fig. 14 is a simplified interface tool 1400 used to modify chart styles, scales, charting intervals etc. This tool is merely an example and should not limit the scope of the claims herein. The present tool has an "OK" button for saving or storing selected chart options. A "cancel" button 1403 is also shown to delete or remove selected chart options. A help button 1405 is shown to identify features of any of the chart options. Numerous chart options 1407 exist. For example, options include, among others, a legend, a value bar, a vertical grid, a horizontal grid, and vertical labels. To select any one of these options, the user clicks onto the box located next to the option or enters the underlined key designating the option. Chart options also include a gallery 1409, either in graph or histogram form. Additionally, the chart can have a maximum vertical scale 1411 such as the 56 for 56 kbits/second. Furthermore, the chart can have a refreshing or updating cycle time 1413. In one aspect, the cycle time can be manually updated. Alternatively, the cycle time can be periodically updated. When using the periodically updating feature, a time interval (e.g., seconds) needs to be specified and entered into a field, as shown.

Fig. 15 is a simplified graphical user interface 1500 for adding or specifying an additional chart according to the present invention. This interface or tool is merely an example and should not limit the scope of the claims herein. This interface allows the user to select the parameters to be monitored on the chart. These parameters include, among others, the computer to be monitored 1507, the object 1509, the counter 1511, and the instance 1514. Depending on the types of parameters being monitored or profiled, specific visual details of the plots or charts are also selected. These details include the plot color 1513, the plot width 1519, the plot style 1517, and others. A counter definition 1515 is also made or selected. Once all the changes have been made or selected, the user can add the changes to be monitored by the tool by pressing or selecting the add button 1501. Alternatively, the user may start over by selecting the cancel button 1503. If the user would like an explanation on any one of the features described in the tool, the user may select either the explain button 1505 or the help button 1506. Of course, this user interface is merely an example and should not be limiting any manner outside the spirit and scope of the claims.

In yet an alternative aspect, the present monitoring or profiling tool has a save feature for storing the chart or plot. In particular, the present tool can save snapshots of measurements to a disk file or the like. As merely an example, the present tool saves snapshots using the following sequence of steps, which should not be construed as limiting:

Go to view/log in the tool to configure a log file;
Add measurements to the file and start and/or stop logging.

Furthermore, the present tool provides congestion, utilization, and performance degradation reports, which make day to day troubleshooting much simpler and serve to justify or validate policy setting decisions. For example, a chronic problem affecting a service through a day period (i.e., 24 hour) can be monitored by a combination of real-time monitoring, which will be described in more detail below, and congestion reports. By monitoring and using the reports, it may be determined that the affected service is not getting its due share of bandwidth, or a limitation exists with the server or in the Internet backbone.

DISTRIBUTED BANDWIDTH MANAGEMENT

In a specific embodiment, the present invention provides techniques for distributed bandwidth management. The present distributed bandwidth management tool can be used to enable and implement end-to-end QoS ("EQOS") and full-cycle Traffic Management ("FTM") in an enterprise network, the Internet, and the like.

Before proceeding to the specific embodiments, however, it may assist the reader to understand the following definitions, which should not be limiting.

1. Distributed policy management ("DPMA"): DPMA is a technique or solution that enables and implements end-to-end QoS (EQOS) and full-cycle Traffic Management (FTM) in an enterprise network and the Internet.
2. Full-Cycle Traffic Management ("FTM"): Traffic management with feedback control for reporting and/or monitoring.

3. End-end Quality of Service ("EQOS"): EQOS enables end-end controls if necessary. In one embodiment, an EQOS Agent is installed on the participating client and server stations. It is used to enable EQOS support in cases where this cannot be done by proxy using routing/switching/fire walling devices.

5

4. Flow Analysis and Measurement Engine ("FAME"): FAME detects and measures traffic by business application and/or transaction. FAME provides components at a client and/or server to accurately detect requests to business applications. It incorporates an application definition language which also specifies measurement methods for the application. FAME can also be incorporated into a gateway product, a firewall, a switch, or router.

10 5. Distributed Bandwidth Broker ("DBB"): Server based distributed bandwidth management engine for allocating traffic in an enterprise network or Internet.

15

6. User Resolution Service ("URS"): URS transparently detects traffic by users who log into a network or server. An enterprise network may choose to implement a variety of techniques to authenticate users, including password based internal directories, remote access directories, firewall directories, tokens or smart cards. 20 URS provides components that run on the client, directory or security server to detect login events and track changes to a users network location (e.g., address). This provides DBMA with the means to enforce policies based on named users/groups, no matter how and where they login from.

25

7. Enterprise Traffic Server ("ETS")/Enterprise Policy Server ("EPS"): ETS/EPS performs traffic analysis and is policy mediation server. ETS/EPS generally maintains the integrity of QoS for all requesting and enforcing components in the DBMA solution. Some functions include:

30

Dynamic traffic and policy analysis; utilizing active monitoring of devices or probing of the network;

Translates policies into dynamic actions that are communicated to enforcement devices via a policy exchange protocol or a standard network management protocol, e.g., SNMP, TELNET;

Provides security to the QoS solution, so that no entity can work around the enterprise policy;

Published statistics to management components; and

Translates external events into changes in policy (event-driven policies).

5

8. Enterprise Policy Manager ("EPM"): EPM creates traffic policy, leveraging existing management directories. EPM allows the creation and validation of business-level enterprise-wide policies relating to traffic management and security, including policies that are created or tuned by events from other programs or network 10 management solutions. EPM transparently leverages existing directory services to access previously defined management objects (e.g., users/groups) and for storing policies.

9. Directory Access and Resolution Service ("DARE"): DARE is generally a directory access and resolution engine to access multiple directories. That is 15 one or more directories can be accessed by way of this resolution engine.

The above definitions are merely intended to assist the reader in understanding some of the terms described herein. They are not intended, in any manner, to limit the scope of the claims. One of ordinary skill in the art would recognize other 20 variations, modifications, and alternatives.

The present distributed bandwidth management embodiment has one of more of these benefits. In one aspect, the present invention provides a business-driven policy management tool for users and/or transactions. The present invention or tool can be implemented into a network without substantial changes to the network infrastructure. 25 The present tool also can provide control at, for example, a user and for a business transaction. The present invention can also be implemented to work with a variety of data rates, e.g., 100 Mbps to 56 Kbps. Furthermore, the present invention can support a full range of traffic controls, including queuing, precise rate control, congestion control, reservation, class-based allocation and prioritization, and others. Still further, the present 30 invention supports a server-based and router/switch based controls, separately or together. In other embodiments, the present invention provides for integrated monitoring, event handling, and event-driven policy settings. These and other benefits are described in more detail below with reference to the Figs.

Fig. 16 is a simplified diagram 1600 of a distributed policy management ("DPMA") system according to an embodiment of the present invention. This diagram is merely an illustration and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives.

- 5 The system 1600 includes a variety of elements. They include a plurality of switches including a workgroup switch 1601, a server switch 1603, a backbone switch 1605, and others. The system also includes routers such as a data center router 1607, a WAN router 1609, and others. The workgroup switch couples to the backbone switch. The workgroup switch also couples to a policy manager 1611, client/server 1613, an IP
- 10 video/phone 1615, and voice gateway 1617, which includes a phone 1619. The WAN router is coupled between the backbone switch and wide area network 1621, e.g., Internet. The switch 1603 couples to application server 1623 and video server 1625. The switch 1603 also connects to data center router 1607, which is connected to the backbone switch. The backbone switch is connected to a variety of elements such as policy services 1627,
- 15 IP call server 1629, DNS/DHCP 1631, NOS authentication 1633, directory services 1635, and other elements. The present system is merely an example and should not limit the scope of the claims herein.

A variety of quality of services ("QoS") agents are distributed throughout the network. One agent is placed in the voice gateway. One agent is placed in the application and video servers. One agent is at the policy services. An agent can also be placed selectively at other locations of the network. Each agent is used to monitor and control bandwidth using one of the techniques described herein as well as others. Each agent is also coupled to the present DPMA tool. Although the system is generally described in terms of one or more agents, the system is not limited to such one or more agents.

In one aspect, DPMA includes, among other elements, the following components. DPMA has FAME, which is used to detect and measure traffic by application and transaction. DPMA also has URS, which transparently detects traffic by user(s) who log in to the network or server, and EPM, which creates traffic policy, leveraging existing management directories. DPMA includes EPS/ETS for traffic analysis and policy mediation server. An EQOS Agent, which enables end-to-end controls if necessary, also is included in the DPMA.

The present system includes a rich set of network services that can be managed using a policy-based approach. As merely an example, policies are rules that govern the behavior of the networking infrastructure in providing services such as QoS, security, and voice/video. Policies are usually stored in a database such as an LDAP compliant directory.

In one embodiment, the present invention provides a technique for policy management. The technique includes a framework for managing a services based network. The technique also includes delivering and implementing business goals by ensuring that the policies that govern the network reflect those goals. The present invention also provides for a "self healing" network, which brings the network back to its normal operating state upon changes. The policy-based, directory-enabled approach allows services to be provided in accordance with user and application requirements and to provide a much higher degree of automation in the management of the network. This automation not only makes management less labor intensive, it also improves the timeliness and quality of network management – with the ultimate goal being the self-correcting, self-healing network.

Fig. 17 is a simplified diagram 1700 of an intelligent network according to an embodiment of the present invention. This diagram is merely an illustration and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. The intelligent network includes a variety of elements or building blocks, which can be distributed throughout the network. The present network includes a graphical management console through which high level network policies derived from business objectives can be created.

The present network also includes a Meta-Policy Service 1701, which provides event management and communicates with the other elements of the management system such as the directory services 1703, policy-enabled network services, intelligent agents and external systems. In one aspect, directory services include user profiles, network device information, and network policies, which are integrated with address management services for IP address assignment ("DHCP") and name-address resolution ("DNS servers"). The network also includes network services such as QoS 1709, security 1707, accounting and billing 1711, device configuration 1713, and others 1715. Network switching 1705 such as routers, switches, firewalls, and others represent the physical layer of the present network.

The network further includes intelligent agents that monitor and control network traffic, and which have application level intelligence up to layer seven and others. This not only enables them to manage traditional IP services (ftp, telnet, SMTP and so on) but also business applications and multi-media streams that blend voice 5 interactive video and data. These agents may be stand alone or embedded in networking devices such as switches and routers, that can be enabled to intelligently enforce policies such as prioritizing business critical traffic, which gives SAP R/3 traffic priority over PointCastTM stock quote updates, for example. The agent can also reside on application servers, which takes advantage of the intelligence of these end-node devices while off- 10 loading the network devices themselves. This takes advantage of the distributed intelligence of the network and the hosts connected to it.

Given the above background, the present invention includes other aspects of the network elements, which are described below.

The meta-policy service can distribute policies to intelligent agents, 15 maintains network state information, correlates dynamic events, and performs other functions. The service can also take corrective action in accordance with pre-determined network policies such as the ones described herein. For example, such an action might entail dynamically re-allocating bandwidth based on network response time for a critical 20 application, or changing the security permissions of a given user based on his or her move to a different department within a company.

25 Directory services is used to implement policy-based management. The directory services maintain user profiles, network device information as well as network policies. Directory services can also be integrated with address assignment ("DHCP") and address resolution ("DNS") servers. This integration can help to automate more of the administration of the network and to make it simpler to create network policies that are based on higher level objects such as users, groups and organizational units – rather 30 than low level entities such as IP or MAC addresses.

QoS is a service within the intelligent network, and is dependent on a policy management framework. In a dynamic environment of an IP

- network, for example, the ability to ensure that user and application requirements for throughput and response time are met in a timely fashion is important in some cases. It makes it possible for service providers to provide differentiated services. And it makes it possible for IT managers to ensure that business critical applications will not be overwhelmed by competing, low priority application traffic. Similarly, it makes it possible for applications with widely varying latency requirements – voice and data, for example – to co-exist on the same network.
- 10 Security services include functions such as access control, authentication, authorization, and encryption.
- 15 Device configuration is one of the more complex tasks associated with the management of the network, in that it is a highly device specific task. This function tends therefore to be the province of each individual hardware vendor. But it lends itself to a policy management approach in helping to automate tedious functions such as software updates on routers and switches.
- 20 Accounting and billing services go hand in hand with the differentiated services approach to managing the network. Since the value of information varies greatly by its timeliness, content, source and destination, it makes sense to charge for service on the basis of the value delivered.
- 25 Network devices include, for example, switches, routers and firewalls. Increasingly, these will tend to be intelligent devices with embedded intelligent agents, so that they become active participants in the policy management framework.
- 30 Fig. 18 is a simplified diagram 1800 of an intelligent network according to an alternative embodiment of the present invention. This diagram is merely an illustration and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Similar to the

Internet itself, the present intelligent management system can be implemented in a decentralized fashion. Decentralization provides a variety of benefits of more flexibility, more resilience, more scalability and a greater adaptability in meeting future and/or ongoing needs. The present system includes a policy manager 1801 that oversees a 5 network. The policy manager includes directory access and a resolution engine 1803. The directory access and resolution engine interface with policy-driven services 1805. These engines monitor and control enablement and enforcement agents 1813. Additionally, the directory access interfaces with management stations network services 10 1811 through SNMP. Additionally, directory access interface switch business systems 1809 via application protocol interface ("API"). The directory access also interfaces with directories and DNS/DHCP address management. Details of the above elements are described in more detail below.

15 1) The Policy Management interface or policy manager: The manager is implemented in, for example, a Web based, graphical user interface. The interface is preferably on an MIS managers desk, for example. The interface can be used for a variety of functions such as to profile the network, create policies and to view their results. In some embodiments, the manager can be implemented at one of many locations. That is, it provides a location independent user interface along the network. 20 Additionally, the manager can be a Java-based browser that can be implemented anywhere on the network, local, or remote.

25 As merely an example, Fig. 19 illustrates a screen 1900 or graphical user interface (GUI) from the manager, which serves to illustrate the look and feel of a policy management interface. The present Fig. is merely an illustration and should not limit the scope of the claims herein. The user interface has been configured as rows and columns, where the rows represent a traffic class or category. Each column represents one of many features for each class or category of traffic. The feature can be, for example, a rule 1901, 30 a sender 1903, a receiver 1905, a service 1907, time 1909, bandwidth allocated 1911, priority 1913, and admissions 1915. Through this GUI, the present invention uses policies to define both monitoring and control actions. These rules incorporate: a traffic class which defines a flow or set of flows including source, destination, application and

file type. Traffic classes can incorporate users and groups – which may be defined in repositories such as:

- NTTM domains or Novell's NDS;
- A timeframe;
- 5 A QoS policy (e.g., a bandwidth reservation, a limit or guarantee, a priority level).

10 The type of signaling or enforcement (e.g., WFQ, RSVP, IP Precedence) used by that agent. These can be either manually configured or, through integration with network management platforms, they can be discovered. The GUI also allows configuring event-triggered actions, such as invoking a QoS control policy when network response time detected by an intelligent agent falls below the required minimum threshold. Alarms and notifications can also be specified, in order to determine which events will trigger an alarm, at what threshold, and in what form e.g. email notification, pager message, SNMP 15 trap, log entry and so on.

2) The Meta-Policy Service performs multiple functions.

20 The service provides a meta-directory function, through interfacing to one or more directory services via a common API, using either LDAP or proprietary protocols such as the one used by Novell's NDS. It also interfaces with address management services such as DNS and DHCP (and is aware of DHCP leases). These functions may be critical for interfacing with a variety of directory services (e.g., LDAP and non-LDAP) and for ensuring that it is compliant in environments where no directories exist or 25 where traffic classes will be defined using entities other than a directory entity – such as a DNS name, for example. The service receives dynamic network event information from external systems such as RMON probes and business applications, and can in turn distribute network event information to other systems. And the service correlates and evaluates 30 events, in order to provide intelligent event handling for the various policy-driven services.

3) Policy Services evaluate and interpret policy requests. In responding to policy requests, PDP, these "policy decision engines" take several variables into account: pre-configured business rules, the nature of the request, dynamic events, network topology and the state of the network (e.g. utilization). Multiple instances of any Policy Service can be deployed in order to improve scalability. Each will handle the devices and agents within its administrative domain.

5
10 Note: Other services such as a RADIUS server and an H.323 gatekeeper could also be considered to be policy servers.

15
4) Policy Transaction protocols: Such protocols can be an essential piece of the policy management framework, and include protocols such as RADIUS, COPS and DIAMETER. Initially, the Policy Server will often support COPS RSVP v1 clients and agents (via COPS extensions)

20
5) Policy Proxy: This is a software module which can "push" policies (i.e. configure enforcement policies) on devices which are not "policy aware". A policy aware device is one that is capable of contacting a Policy Decision Engine via a policy Transaction Protocol. The present device can support QoS configuration of Cisco and Bay network devices. The communication mechanism will be SNMP SET's and or HTTP and or CLI and/or Telnet.

25
6) Intelligent agents are of at least two kinds (which may run on the same device):

30
Policy Enforcement Points: These reside in the data path and have enforcement capabilities such as blocking/dropping/ queuing/modifying packets as they flow through the enforcement point. If a policy enforcement node does not have sufficient context to make an enforcement decision, it may refer to another node, a policy decision engine, via a policy transaction protocol such as COPS. Examples of such agents are

5 routers, switches, firewall agents and others running on an NT or UNIX servers. These agents may be enforcing a policy based on a policy that has been statically configured, e.g., give highest priority to any SAP R/3 packets or simply by honoring the priority that is signaled by RSVP or Type of Service (ToS) bits set in the incoming packet.

10 Policy Enablement Points or Proxies: These also commonly reside on the data path and are capable of QoS signaling and can do so on behalf of legacy applications which may not be QoS aware. For example, the present agent might reside on an NT or UNIX application server and enable IP Precedence support by setting the TOS bit in the IP header of packets generated by a critical application such as an SQL database application. Downstream enforcement points would then enforce this TOS-defined priority level using, for example, the multi-level queuing capabilities of a router. This enablement function implemented on the end node can alleviate the processing burden that is otherwise imposed on network devices such as routers and can eliminate the need for expensive upgrades to the router infrastructure.

15

20 Depending upon the embodiment, one or more of these advantages may be present. In one aspect, the present invention provides an open, standards-compliant, software-based application. In particular, the present tool is software-based and has an open architecture that gives it flexibility. This approach enables the present tool to fit into a wide variety of enterprise and service provider network environments, to be easily 25 integrated with network devices provided by OEM partners, and to inter-operate with traffic management systems provided by other vendors.

30 In an alternative aspect, the present invention provides a very simple tool by way of its user interfaces. The present tool also is a Java-based Web interface gives it the location independence required for "manage from anywhere" administration. Another aspect to simplifying the network manager's life is reducing the number of data repositories that have to be maintained and synchronized. The present tool's "meta directory"-like integration with NT Domains, LDAP directories and DNS/DHCP servers

means that implementing it does not require duplicating yet another data store – and in fact, the use of any directory service at all is optional.

In one or more embodiments, the present policy management framework provides a policy-based, directory-enabled traffic management, as well as other features.

- 5 The present invention also provides a fully distributed traffic management system in other embodiments. In one or more aspects, the present system is designed to be “open,” standards-compliant, scalable and robust. The present invention can also extended to support not only QoS but also the full range of network services that are elements of the intelligent network.

10

EXAMPLE

Background

As merely an example, a sample network is configured to carry out aspect of the present invention. This is merely an example and should not limit the scope of the 15 claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. In the present example, the hypothetical company is called “Company.” The following parameters must be met by the network in the Company. Company’s network includes a local area internal network and an external wide area to corporate offices as well as the Internet.

20

Company has “USER Max” which is a road warrior. USER Max must be guaranteed 128 Kb of bandwidth to access the corporate servers from all remote offices over leased lines, by dialing into the corporate network, or over the Internet using a VPN solution.

25

A mission critical application is called “CashReg.” CashReg must be guaranteed 30% of all leased line links from remote offices. In addition, each user using this application should be guaranteed 56 Kb, subject to a limit of 80% for the application as a whole.

30

Company’s Web-Site must be guaranteed 20% of the Internet access link with high-priority.

The present bandwidth management tool enforces the above parameters in the following manner.

For User Max: DPMA URS component detects any logins by User Max
5 and when User Max starts accessing any service, the WAN router, dial-in router or firewall serving User Max will be signaled by DPMA Policy Server to allocate at least 128K bandwidth for User Max against other competing traffic.

For Application "CashReg": DPMA FAME component detects any access
10 to CashReg and the DPMA Policy Server keeps track of the application as a whole, and communicates the bandwidth requirement to all the appropriate WAN routers or DPMA application server components. The per-user 56 Kb minimum is also enforced, but if User Max uses CashReg, he will get 128Kb.

15 For the Web Server: High priority Internet service is communicated by the Policy Server to Web server EQOS Agent. The Agent sets IP Type of Service for all traffic from the Web server. In addition, the Policy Server communicates the bandwidth requirement to the Internet router, which supports this function. If this function is not supported, the Server EQOS Agent will assume control of bandwidth as well.

20

In the foregoing specification, the invention has been described with reference to specific exemplary embodiments thereof. Many changes or modifications are readily envisioned. For example, the present invention can be applied to manage a variety of TCP/IP network traffic types for the Internet and Intranet. Further, the
25 techniques can also be applied to Novell SPX, Xerox XNS or any protocol with a similar 'flow-control' design that utilizes windows and acknowledgment signals (similar to ACK). Alternative embodiments of the present invention can also be applied to a 'legacy' private WAN running IP as well as native Novell protocols if there is a need. (e.g., file server and client/server traffic). Further, embodiments of the present invention can include
30 monitoring, billing, and reporting features, thus allowing for enhanced client billing and internal cost accounting of network usage.

Furthermore, the above descriptions have been described in terms of bandwidth management generally and a distributed bandwidth management embodiment.

- It would be recognized, however, that aspects of the general bandwidth management can be combined with the distributed embodiment. These embodiments also can be separated and then recombined with other features. Accordingly, the invention should not be limited to the description in the specific embodiments described. These techniques are
- 5 preferably implemented within a firewall platform to solve the provide the following benefits: bi-directional bandwidth management of network links carrying TCP traffic; reactive (short-time scale) and proactive (long time scale) control mechanisms; and gateway (local) and end-end (global) techniques for bandwidth control. This solution reduces their contribution to congestion in the Internet; and operation in a present day
- 10 heterogeneous wide area networks, such as the Internet, without requiring any client, server or router changes.

The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. It will, however, be evident that various modifications and changes may be made thereunto without departing from the broader

15 spirit and scope of the invention as set forth in the claims.

WHAT IS CLAIMED IS:

1. A graphical user interface for monitoring a flow of information
2 coupled to a network of computers, said graphical user interface comprising:
 - 3 a display comprising at least a first portion and a second portion, said first
4 portion comprising a graphical chart representing said flow of information, said second
5 portion comprising text information describing said flow of information;
 - 6 whereupon said display also outputs a network location from one of a
7 plurality of selected locations for said flow of information.
- 1 2. The interface of claim 1 wherein said graphical chart comprises
bandwidth consumption.
- 1 2. The interface of claim 2 wherein said bandwidth consumption is a
plot of bandwidth consumed against time for said network location.
- 1 2. The interface of claim 2 wherein said bandwidth consumption is a
plurality of plots, each of said plots representing consumed bandwidth against time.
- 1 2. The interface of claim 2 wherein said flow of information
comprises one of a plurality of traffic classes.
- 1 2. The interface of claim 1 wherein graphical chart comprises a plot
of failure rates against time
- 1 2. The interface of claim 1 wherein said graphical chart comprises a
plot of delay rates against time.
- 1 2. The interface of claim 1 wherein said display is outputted on a
computer monitor.
- 1 2. The interface of claim 1 wherein said display is a real-time display
of a portion of said flow of information.
- 1 2. The interface of claim 1 wherein said network location is selected
from a graphical representation or text.

1 11. An intelligent computer network system, said computer network
2 system comprising a real-time distributed bandwidth profiling tool, said real-time
3 bandwidth profiling tool comprising a graphical user interface on a monitor, said
4 graphical user interface comprising at least a first portion and a second portion, said first
5 portion comprising a graphical chart representing said flow of information, said second
6 portion comprising text information describing said flow of information, whereupon said
7 flow of information is derived from one of a plurality of selected network locations
8 distributed over said computer network.

1 12. A method for monitoring a flow of information coupled to a
2 network of computers, method comprising:

3 providing a display comprising at least a first portion and a second portion,
4 said first portion comprising a graphical chart representing said flow of information, said
5 second portion comprising text information describing said flow of information;

6 whereupon said display also outputs a network location from one of a
7 plurality of selected locations for said flow of information.

1 13. The method of claim 12 wherein said graphical chart comprises
2 bandwidth consumption.

1 14. The method of claim 13 wherein said bandwidth consumption is a
2 plot of bandwidth consumed against time for said network location.

1 15. The method of claim 13 wherein said bandwidth consumption is a
2 plurality of plots, each of said plots representing consumed bandwidth against time.

1 16. The method of claim 13 wherein said flow of information
2 comprises one of a plurality of traffic classes.


1 17. The method of claim 12 wherein graphical chart comprises a plot
2 of failure rates against time

1 18. The method of claim 12 wherein said graphical chart comprises a
2 plot of delay rates against time.

1 19. The method of claim 12 wherein said display is outputted on a
2 computer monitor.


1 20. The method of claim 12 wherein said display is a real-time display
2 of a portion of said flow of information.

1 21. The method of claim 12 wherein said network location is selected
2 from a graphical representation or text.



SCHEMATIC OF QOS SYSTEM DEPLOYED FOR INTERNET ACCESS
AND A PRIVATE WAN

FIG. 1






FIG. 3




FIG. 4






FIG. 5



↙ 600

FIG. 6



↙ 700

FIG. 7




FIG. 8




FIG. 9A




Fig. 9

~ 1000




FIG. 10

~ 1100




FIG. 11




FIG. 12




FIG. 13

21400




FIG. 14

21500





FIG. 15







FIG. 17



15 / 15




FIG. 19

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US98/27774

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : H04C 12/66; H04J 3/14
US CL : 370/401, 252, 253, 468; 709/223, 224, 225, 226

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 370/401, 252, 253, 468; 709/223, 224, 225, 226

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X.P	US 5,825,775 A (CHIN et al) 20 October 1998, See Abstract, and Fig. 3.	1-21
X.P	US 5,867,483 A (ENNIS, Jr. et al) 02 February 1999, See Abstract, Fig 1 and Fig 13-16.	1-21
X.P	US 5,819,028 A (MANGHIRMALANI et al) 06 October 1998, See Abstract and Fig 1-20.	1-21

 Further documents are listed in the continuation of Box C. See patent family annex.

A	Special categories of cited documents	*T*	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
B	document defining the general state of the art which is not considered to be of particular relevance	*N*	document of particular relevance, the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
C	earlier document published on or after the international filing date	*Y*	document of particular relevance, the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
D	document which may throw doubts on priority claims(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*E*	document member of the same patent family
F	document referring to an oral disclosure, use, exhibition or other means		
P	document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search

10 FEBRUARY 2000

Date of mailing of the international search report

23 FEB 2000

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

STEVEN NGUYEN

Telephone No. (703) 308-8848